

IN THE CLAIMS

Please amend the claims as follows:

1. (Currently Amended) A method of accessing shared memory in a computer system having a plurality of nodes, including a first node, wherein each node includes a processor, a Translation Look-aside Buffer (TLB) associated with the processor and local memory wherein the local memory of each node includes a Remote Translation Table (RTT), the method comprising:

distributing an application across the plurality of nodes including the first node;

building an application virtual address space in each of the plurality of nodes, wherein building an application virtual address space includes:

building a local virtual address space for the application in each of the plurality of nodes, wherein[[::]] the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node [[;]] and wherein the virtual address generated by the application executing on the node includes a node number of the node; and

exporting the local virtual address space for each local node to a Remote Translation Table (RTT) associated with that node the RTTs in each of the plurality of nodes, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation; and

performing a virtual memory reference to a physical memory location, wherein performing the virtual memory reference includes:

translating the virtual memory reference into a physical address in the application virtual address space using the RTT on the local node, if the node number is not the local node number and remote translation is enabled, wherein performing a translating the virtual memory reference to a memory location into a physical address in the application virtual address space includes translating the node number of the application virtual address into a node address associated with the first node and translating bits of the application virtual address into a physical page address for the first node using the RTT

associated with the first node; and

translating the virtual memory reference into a physical address in local memory for the first node using the RTT on the first node, if the node number is not the local node number and remote translation is not enabled, wherein translating the virtual memory reference includes sending the bits of the application virtual address to the first node.

2. (Currently Amended) The method of claim 1, wherein the local address space is read from a the Translation Look-aside Buffer (TLB) associated with the node.
3. (Original) The method of claim 1, wherein building an application virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.
4. (Currently Amended) A system comprising:
a plurality of nodes, each node including:
one or more processors having a Translation Look-aside Buffer (TLB);
a memory; and
a memory controller operatively coupled to the memory and the one or more processors, wherein the memory controller includes a Remote Translation Table (RTT), wherein the RTT translates a virtual address received as part of a memory request received from another node into a memory request with physical addresses into the memory on the node associated with the RTT; further wherein the RTT is initialized upon the start of a process associated with an application by building virtual to physical address translations for local virtual address space in the node corresponding to the application, wherein a each virtual address includes a node number of the node, and by exporting the virtual to physical address translations for the local virtual address space from the node to the Remote Translation Table (RTT) RTT in each node associated with that node through the application, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective local node and requesting that the operating system enable remote

translation[[.]];

wherein the TLB translates a virtual address reference received from the local node into a physical address in the memory for the local node, if the node number for the virtual address is the local node number;

wherein the RTT translates the virtual address received from the local node into a physical address in the memory for a remote node associated with the local node, if the node number for the virtual address is not the local node number and remote translation is enabled; and

wherein the RTT sends the virtual address received from the local node to another node and has the RTT on that node translate the virtual address into a physical addresses in the memory for that node, if the node number for the virtual address is not the local node number and remote translation is not enabled.

5. (Original) The system of claim 4, wherein each of the plurality of nodes executes a synchronization operation that causes at least some of the plurality of nodes to wait for all of the plurality of nodes to complete exporting the virtual to physical address translations to their respective Remote Translation Tables.

6. (Currently Amended) A device-readable medium having instructions thereon that, when executed on a properly programmed information-processing device having a plurality of nodes, including a first node, each node having one or more processors with a Translation Look-aside Buffer (TLB), a memory, and a memory controller and coupled to the memory and the one or more processors, causes the information-processing device to perform a method comprising:

distributing an application across the plurality of nodes including the first node;

building an application virtual address space in each of the plurality of nodes, wherein building an application virtual address space includes:

building a local virtual address space for the application in each of the plurality of nodes, wherein[[.]] the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node [[.]] and wherein the virtual address generated by the application executing on the

node includes a node number of the node; and

exporting the local virtual address space for each local node to a Remote Translation Table (RTT) associated with that node the RTTs in each of the plurality of nodes, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation; and

performing a virtual memory reference to a physical memory location, wherein performing the virtual memory reference includes:

translating the virtual memory reference into a physical address in the application virtual address space using the RTT on the local node, if the node number is not the local node number and remote translation is enabled, wherein performing a translating the virtual memory reference to a memory location into a physical address in the application virtual address space includes translating the node number of the application virtual address into a node address associated with the first node and translating bits of the application virtual address into a physical page address for the first node using the RTT associated with the first node; and

translating the virtual memory reference into a physical address in local memory for the first node using the RTT on the first node, if the node number is not the local node number and remote translation is not enabled, wherein translating the virtual memory reference includes sending the bits of the application virtual address to the first node.

7. (Original) The device-readable medium of claim 6, wherein building a local virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes complete exporting their respective address space.

8. (Currently Amended) The device-readable medium of claim 6, wherein the local address space is read from a the Translation Look-aside Buffer (TLB) associated with the node.

9. (Currently Amended) A multinode system for implementing remote address translation, the system comprising:

a plurality of nodes, including a first node, each of the plurality of nodes including:

one or more processors with a Translation Look-aside Buffer (TLB);[[,]]

a memory[[,]]; and

a memory controller operatively coupled to the memory and the one or more processors, wherein the memory controller includes a Remote Translation Table (RTT);

means for distributing an application across the plurality of nodes including the first node;

means for building an application virtual address space in each of the plurality of nodes, wherein the means for building an application virtual address space includes:

means for building a local virtual address space for the application in each of the plurality of nodes, wherein[[,]] the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node [[,]] and, wherein the virtual address generated by the application executing on the node includes a node number of the node; and

means for exporting the local virtual address space for each local node to a Remote Translation Table (RTT) the RTT in each of each the plurality of nodes, associated with that node, wherein the means for exporting includes means, within each node, for requesting that the operating system load the RTT from the local address space of its respective node and means for requesting that the operating system enable remote translation; and

means for performing a virtual memory reference to a physical memory location, wherein means for performing the virtual memory reference includes:

means for translating the virtual memory reference into a physical address in local memory using the TLB on the local node, if the node number is the local node number;

means for translating the virtual memory reference into a physical address in the application virtual address space for the local node using the RTT on the local node, if the node number is not the local node number and remote translation is enabled, wherein performing a means for translating the virtual memory reference to a memory location into a physical memory in the application virtual address space includes[[,]] means for translating the node number of the application virtual address into a node address

associated with the first node[.,.] and means for translating bits of the application virtual address into the physical page address for the first node using the RTT associated with the first node[.,.]; and

means for translating the virtual memory reference into a physical address in local memory for the first node using the RTT on the first node, if the node number is not the local node number and remote translation is not enabled, wherein means for translating the virtual memory reference includes means for sending the bits of the application virtual address to the first node.

10. (Original) The multinode system of claim 9, wherein building an application virtual address space further includes a means for performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.

11. (Currently Amended) A multi-node system for implementing remote address translation, the system comprising:

a network;

a source node coupled to the network, wherein the source node includes a first remote-translation table (RTT); and

a remote node coupled to the network, wherein the remote node includes a second RTT; wherein on the remote node the second RTT is built using a first local address space on the source node exported from the source node to the remote node using an operating system call to perform the export;

wherein on the source node the first RTT is built using a second local address space on the remote node exported from the remote node to the source node using the operating system call to perform the export;

wherein the operating system enables remote translation utilizing the first and second RTTs, wherein enabling remote translation utilizing the first and second RTTs includes having the remote node translates a virtual memory address associated with the source node to a physical address on the source node as a function of the second RTT and having the source node

translates a virtual memory address associated with the remote node to a physical address on the remote node as a function of the first RTT; and

wherein both the first and second RTTs include one or more virtual address and each virtual address includes a node number of a remote node that built the virtual address.

12. (Previously Presented) The method of claim 1, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.

13. (Previously Presented) The method of claim 12, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

14. (Previously Presented) The system of claim 4, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.

15. (Previously Presented) The system of claim 14, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

16. (Previously Presented) The device-readable medium of claim 6, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.

17. (Previously Presented) The device-readable medium of claim 16, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

18. (New) The method of claim 1, wherein requesting the operating system enable remote translation handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

19. (New) The system of claim 4, wherein, when remote translation is enabled, the operating system handles requests to changes the application virtual address space configuration on a node-

local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

20. (New) The device-readable medium of claim 6, wherein requesting the operating system enable remote translation handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

21. (New) The system of claim 9, wherein, when remote translation is enabled, the operating system handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

22. (New) The system of claim 11, wherein, when remote translation is enabled, the operating system handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.